Porting with HPE's MPI and Intel OpenMP

Building Applications

To build an MPI/OpenMP hybrid executable using HPE's MPT and Intel's OpenMP libraries, you must compile your code with the **-openmp** flag and link it with the **-lmpi** flag. For example:

```
%module load comp-intel/2020.4.304 mpi-hpe/mpt
%ifort -o your_executable prog.f -openmp -lmpi
```

Running Applications

Here is a sample PBS script for running an MPI/OpenMP application on Pleiades using 3 nodes, with 10 MPI processes on each node and 2 OpenMP threads per MPI process.

```
#PBS -lselect=3:ncpus=20:mpiprocs=10:model=ivy
#PBS -lwalltime=1:00:00

module load comp-intel/2020.4.304 mpi-hpe/mpt
setenv OMP_NUM_THREADS 2

cd $PBS_O_WORKDIR

mpiexec ./your_executable
```

You can specify the number of threads on the PBS resource request line using **ompthreads**. This causes the PBS prologue to set the **OMP_NUM_THREADS** environment variable.

```
#PBS -lselect=3:ncpus=20:mpiprocs=10:ompthreads=2:model=ivy
#PBS -lwalltime=1:00:00

module load comp-intel/2020.4.304 mpi-hpe/mpt

cd $PBS_O_WORKDIR

mpiexec ./your_executable
```

Performance Issues

For pure MPI codes built with HPE's MPT library, you can improve performance on Sandy Bridge, Ivy Bridge, Haswell, and Broadwell nodes by pinning the processes through setting **MPI_DSM_DISTRIBUTE** environment variables to 1 (or true). However, for MPI/OpenMP codes, all the OpenMP threads for the same MPI process have the same process ID. In this case, setting this variable to 1 causes all OpenMP threads to be pinned on the same core and the performance suffers.

It is recommended that you set **MPI_DSM_DISTRIBUTE** to 0 and use **omplace** for pinning instead.

If you use Intel version 10.1.015 or later, you should also set **KMP_AFFINITY** to *disabled* or **OMPLACE_AFFINITY_COMPAT** to ON as Intel's thread affinity interface would interfere with **dplace** and **omplace.**

```
#PBS -lselect=3:ncpus=20:mpiprocs=10:ompthreads=2:model=ivy
#PBS -lwalltime=1:00:00

module load comp-intel/2020.4.304 mpi-hpe/mpt

setenv MPI_DSM_DISTRIBUTE 0
setenv KMP_AFFINITY disabled
```

```
cd $PBS_O_WORKDIR

mpiexec -np 30 omplace ./your_executable
```

Article ID: 104
Last updated: 13 May, 2021
Revision: 25
Porting/Building Code -> Porting to NAS Systems -> Porting with HPE's MPI and Intel OpenMP
https://www.nas.nasa.gov/hecc/support/kb/entry/104/